arXiv:1607.04109v1 [cs.IT] 14 Jul 2016

General Sub-packetized Access-Optimal Regenerating Codes

Katina Kralevska, Danilo Gligoroski, Harald Øverby

***Abstract*—This paper presents a novel construction of $(n, k, d = n-1)$ access-optimal regenerating codes for an arbitrary sub-packetization level α for exact repair of any systematic node. We refer to these codes as general sub-packetized because we provide an algorithm for constructing codes for any α less than or equal to $r^{\lceil \frac{k}{r} \rceil}$ where $\frac{k}{r}$ is not necessarily an integer. This leads to a flexible construction of codes for different code rates compared to existing approaches. We derive the lower and the upper bound of the repair bandwidth. The repair bandwidth depends on the code parameters and α. The repair process of a failed systematic node is linear and highly parallelized, which means that a set of $\lceil \frac{\alpha}{r} \rceil$ symbols is independently repaired first and used along with the accessed data from other nodes to recover the remaining symbols.**

***Index Terms*- Minimum storage regenerating codes, sub-packetization, access-optimal**

I. Introduction

Erasure coding is becoming an attractive technique for data protection since it offers the same level of reliability with significantly less storage overhead compared to replication [1]. Apart from the reliability and the storage overhead, there are other desirable features in a distributed storage system such as low repair bandwidth and access-optimality. Repair bandwidth is the amount of transferred data during a repair process. Access-optimality is achieved when the amount of accessed and transferred data during the repair process is equal.

Dimakis et al. introduced regenerating codes that significantly reduce the repair bandwidth [2]. Under an (n, k, d) regenerating code, a file of M symbols from a finite field $\mathbf{F}_q$ is divided into k fragments, each of size $\alpha = \frac{M}{k}$ symbols, which are further encoded into n fragments using an (n, k) MDS (Maximum-Distance Separable) code. The parameter α, termed as a sub-packetization level of the code, represents the minimum dimension over which operations are performed. The data from a failed node is recovered by transferring β symbols from each d non-failed nodes. Thus, the repair bandwidth γ is equal to $d\beta$ where $\alpha \leq d\beta \ll M$. Dimakis et al. [2] showed the existence of minimum storage regenerating (MSR) codes that attain the minimum storage point of the optimal tradeoff curve between the storage and the repair bandwidth, i.e.,

$$(\alpha_{MSR}, \gamma^{min}_{MSR}) = (\frac{M}{k}, \frac{M}{k}\frac{n-1}{n-k}). \quad (1)$$

The repair bandwidth is minimized when all $d = n-1$ non-failed nodes transmit a fraction of $1/r$ of the stored data.

Several exact repair MSR codes, which are characterized by the repaired data being exactly the same as the lost data, have been suggested. Tamo et al. proposed optimal MSR codes for α equal to r^k known as zigzag codes [3]. Furthermore, they showed that α of access-optimal MSR codes for repair of any systematic node is $r^{\frac{k}{r}}$ [4]. The codes presented in [5] and [6] meet this condition. An essential condition for the code construction in [6] is that $m = \frac{k}{r}$ has to be an integer $m \geq 1$ where k is set to rm and α to r^m. Wang et al. constructed codes that optimally repair any systematic or parity node for α equal to r^{k+1} [7]. High-rate MSR codes with a polynomial sub-packetization level are proposed in [8]. However, our work focuses only on code constructions for optimal repair of any systematic node.

MSR codes are optimal in terms of storage, reliability and repair bandwidth, but not I/O. Implementing MSR codes with a sub-packetization level of $r^{\frac{k}{r}}$ may not be practical for storage systems that serve applications with a large number of user requests or perform intensive computations. Thus, having an algorithm for constructing MSR codes for any combination of n, k and α that are simultaneously optimal in terms of storage, reliability, repair bandwidth and I/O is an important problem that is solved in this work.

Our Contribution: This paper presents a novel construction of $(n, k, d = n-1)$ access-optimal regenerating codes for an arbitrary sub-packetization level for exact repair of any systematic node. The codes have the following properties: 1. MDS; 2. Systematic; 3. Flexible sub-packetization level; 4. Minimum repair bandwidth for every α including the lower bound (1) when α is $r^{\lceil \frac{k}{r} \rceil}$; 5. Access-optimality; 6. Fast decoding. To the best of the authors' knowledge, these are the first code constructions for an arbitrary α. Motivated by the code construction in [6], we construct general codes where $\frac{k}{r}$ does not need to be an integer and α is not exclusively equal to $r^{\frac{k}{r}}$. For instance, the code $(14, 10)$ that is deployed in the data-warehouse cluster of Facebook [9] is out of the scope of applicability with the current proposals in [4]–[6], because $\frac{k}{r} = 2.5$ is a non-integer. However, the presented algorithm constructs an $(14, 10, 13)$ code that reduces the repair bandwidth for any systematic node by 67.5% when α is $r^{\lceil \frac{k}{r} \rceil} = 64$ compared to an $(14, 10)$ RS code. The repair process is linear and highly parallelized.

II. A General $(n, k, d = n-1)$ Code Construction

Consider a file of size $M = k\alpha$ symbols from a finite field $\mathbf{F}_q$ stored in k systematic nodes d_j of capacity α symbols.

We define a systematic MDS code in the following way: The basic data structure component is an index array of size $\alpha \times k$

where $\alpha \leq r^{\lceil \frac{k}{r} \rceil}$ and $n = k + r$, $P = ((i,j))_{\alpha \times k}$. We use r such index arrays $P_1, \ldots, P_r$. The elements $p_{i,1}$, $i = 1, \ldots, \alpha$, in p_1 are a linear combination only of the symbols with indexes present in the rows of P_1. In the initialization phase, additional $\lceil \frac{k}{r} \rceil$ columns with pairs $(0,0)$ are added to $P_2, \ldots, P_r$. The goal of the algorithm is to replace those zero pairs with concrete (i,j) pairs so that the code is access-optimal for a given sub-packetization level α. The value of α determines two phases of the algorithm. In the first phase, the indexes (i,j) that replace the $(0,0)$ pairs are chosen such that both Condition 1 and Condition 2 are satisfied (defined further in this section). The first phase starts with a granulation level parameter called run that is initialized with the value $\lceil \frac{\alpha}{r} \rceil$. This parameter affects how the indexes (i,j) are chosen and with every round the granulation level decreases by a factor r. Once the granulation level becomes equal to 1 and there are still $(0,0)$ pairs that have to get some value (i,j), the second phase starts where the remaining indexes are chosen such that only Condition 2 is satisfied.

A high level description of the proposed algorithm is given in Alg. 1, while a detailed one in Alg. 2.

Algorithm 1 High level description of an algorithm for generating general sub-packetized, access-optimal regenerating codes

1: Initialize the index arrays $P_1, \ldots, P_r$;
2: # Phase 1
3: Set the granulation level $run \leftarrow \lceil \frac{\alpha}{r} \rceil$
4: **repeat**
5: Replace $(0,0)$ pairs with indexes (i,j) such that both Condition 1 and Condition 2 are satisfied;
6: Decrease the granulation level run by a factor r.
7: **until** the granulation level $run > 1$
8: # Phase 2
9: If there are still $(0,0)$ indexes that have to get some value (i,j), choose them such that only Condition 2 is satisfied;
10: Return the index arrays $P_1, \ldots, P_r$;

Once the index arrays $P_1, \ldots, P_r$ are determined, the symbols $p_{i,l}$ in the parity nodes, $1 \leq i \leq \alpha$ and $1 \leq l \leq r$, are generated as a combination of the elements a_{j_1,j_2} where the pair (j_1, j_2) is in the i-th row of the index array P_l, i.e.,

$$p_{i,l} = \sum c_{l,i,j} a_{j_1,j_2}. \tag{2}$$

The linear relations have to guarantee an MDS code, i.e., to guarantee that the entire information can be recovered from any k nodes (systematic or parity). We use the following terms and variables:

- The set $Nodes = \{d_1, \ldots, d_k\}$ of k systematic nodes is partitioned in $\lceil \frac{k}{r} \rceil$ disjunctive subsets $J_1, J_2, \ldots, J_{\lceil \frac{k}{r} \rceil}$ where $|J_\nu| = r$ (if r does not divide k then $J_{\lceil \frac{k}{r} \rceil}$ has $k \mod r$ elements) and $Nodes = \cup_{\nu=1}^{\lceil \frac{k}{r} \rceil} J_\nu$. In general, this partitioning can be any random permutation of k nodes. Without loss of generality we use the natural ordering as follows: $J_1 = \{d_1, \ldots, d_r\}$, $J_2 = \{d_{r+1}, \ldots, d_{2r}\}$, $\ldots$, $J_{\lceil \frac{k}{r} \rceil} = \{d_{\lfloor \frac{k}{r} \rfloor \times r+1}, \ldots, d_k\}$.

Algorithm 2 Algorithm to generate the index arrays
Input: n, k, α;
Output: Index arrays $P_1, \ldots, P_r$.

1: **Initialization:** $P_1, \ldots, P_r$ are initialized as index arrays $P = ((i,j))_{\alpha \times k}$;
2: Append $\lceil \frac{k}{r} \rceil$ columns to $P_2, \ldots, P_r$ all initialized to $(0,0)$;
3: Set $portion \leftarrow \lceil \frac{\alpha}{r} \rceil$;
4: Set $ValidPartitions \leftarrow \emptyset$;
5: Set $j \leftarrow 0$;
6: # Phase 1
7: **repeat**
8: Set $j \leftarrow j + 1$;
9: Set $\nu \leftarrow \lceil \frac{j}{r} \rceil$;
10: Set $run \leftarrow \lceil \frac{\alpha}{r^\nu} \rceil$;
11: Set $step \leftarrow \lceil \frac{\alpha}{r} \rceil - run$;
12: $\mathcal{D}_{d_j} = ValidPartitioning(ValidPartitions, k, r, portion, run, step, J_\nu)$;
13: Set $ValidPartitions = ValidPartitions \cup \mathcal{D}_{d_j}$;
14: Determine one $D_{\rho,d_j} \in \mathcal{D}_{d_j}$ such that its elements correspond to row indexes in the $(k+\nu)$-th column in one of the arrays $P_2, \ldots, P_r$, that are all zero pairs $(0,0)$;
15: The indexes in D_{ρ,d_j} are the row positions where the pairs (i,j) with indexes $i \in \mathcal{D} \setminus D_{\rho,d_j}$ are assigned in the $(k+\nu)$-th column of $P_2, \ldots, P_r$;
16: **until** $(run > 1)$ AND $(j \neq 0 \mod r)$
17: # Phase 2
18: **while** $j < k$ **do**
19: Set $j \leftarrow j + 1$;
20: Set $\nu \leftarrow \lceil \frac{j}{r} \rceil$;
21: Set $run \leftarrow 0$;
22: $\mathcal{D}_{d_j} = ValidPartitioning(ValidPartitions, k, r, portion, run, step, J_\nu)$;
23: Set $ValidPartitions = ValidPartitions \cup \mathcal{D}_{d_j}$;
24: Determine one $D_{\rho,d_j} \in \mathcal{D}_{d_j}$ such that its elements correspond to row indexes in the $(k+\nu)$-th column in one of the arrays $P_2, \ldots, P_r$, that are all zero pairs $(0,0)$;
25: The indexes in D_{ρ,d_j} are the row positions where the pairs (i,j) with indexes $i \in \mathcal{D} \setminus D_{\rho,d_j}$ are assigned in the $(k+\nu)$-th column of $P_2, \ldots, P_r$;
26: **end while**
27: Return $P_1, \ldots, P_r$.

- Each node d_j consists of an indexed set of α symbols $\{a_{1,j}, a_{2,j}, \ldots, a_{\alpha,j}\}$.
- $portion = \lceil \frac{\alpha}{r} \rceil$: The set of all symbols in d_j is partitioned in disjunctive subsets where at least one subset has $portion$ number of elements.
- $run = \lceil \frac{\alpha}{r^\nu} \rceil$, for values of $\nu \in \{1, \ldots, \lceil \frac{k}{r} \rceil\}$.
- $step = \lceil \frac{\alpha}{r} \rceil - run$: For the subsequent $(k+\nu)$-th column, where $\nu \in \{1, \ldots, \lceil \frac{k}{r} \rceil\}$, the scheduling of the indexes corresponding to the nodes in J_ν is done in subsets of indexes from a valid partitioning.
- *A valid partitioning* $\mathcal{D}_{d_j} = \{D_{1,d_j}, \ldots, D_{r,d_j}\}$ of a set of indexes $D = \{1, \ldots, \alpha\}$, where the $i-$th symbol in d_j is indexed by i in D, is a partitioning in r disjunctive subsets $D_{d_j} = \cup_{\rho=1}^{r} D_{\rho,d_j}$. If r divides α, then the valid partitioning for all nodes in J_ν is equal. If r does not divide α, then the valid partitioning has to contain at least one subset D_{ρ,d_j} with $portion$ elements that correspond to the row indexes in the $(k+\nu)$-th column in one of the

arrays $P_2, \ldots, P_r$ that are all zero pairs.

- *Condition 1*: At least one subset D_{ρ,d_j} has $portion$ elements with runs of run consecutive elements separated with a distance between the indexes equal to $step$. The elements of that subset correspond to the row indexes in the $(k+\nu)$-th column in one of the arrays $P_2, \ldots, P_r$ that are all zero pairs. The distance between two elements in one node is computed in a cyclical manner such that the distance between the elements $a_{\alpha-1}$ and a_2 is 2.
- *Condition 2*: A necessary condition for the valid partitioning to achieve the lowest possible repair bandwidth is $D_{d_{j_1}} = D_{d_{j_2}}$ for all d_{j_1} and d_{j_2} in J_ν and $D_{\rho,d_{j_1}} \neq D_{\rho,d_{j_2}}$ for all d_{j_1} and d_{j_2} systematic nodes in the system. If $portion$ divides α, then D_{ρ,d_j} for all d_j in J_ν are disjunctive, i.e., $D = \cup_{j=1}^{r} D_{\rho,d_j} = \{1, \ldots, \alpha\}$.

Algorithm 3 $ValidPartitioning$
Input:$ValidPartitions, k, r, portion, run, step, J_\nu$;
Output: $\mathcal{D}_{d_j} = \{D_{1,d_j}, \ldots, D_{r,d_j}\}$.

Set $D = \{1, \ldots, \alpha\}$;
if $run \neq 0$ **then**
    Find $\mathcal{D}_{d_j}$ that satisfies Condition 1 and Condition 2;
else
    Find $\mathcal{D}_{d_j}$ that satisfies Condition 2;
end if
Return $\mathcal{D}_{d_j}$;

Alg. 4 shows the repair of a systematic node where the systematic and the parity nodes are global variables. A set of $\lceil \frac{\alpha}{r} \rceil$ symbols is accessed and transferred from each $n-1$ non-failed nodes. If $\alpha \neq r^{\lceil \frac{k}{r} \rceil}$, then additional elements may be required as described in Step 4. Note that a specific element is transferred just once and stored in a buffer. For every subsequent use of that element, the element is read from the buffer and further transfer operation is not required. The repair process is highly parallelized, because a set of $\lceil \frac{\alpha}{r} \rceil$ symbols is independently and in parallel repaired in Step 2 and then the remaining symbols are recovered in parallel in Step 5.

Algorithm 4 Repair of a systematic node d_l
Input: l;
Output: d_l.

1: Access and transfer $(k-1)\lceil \frac{\alpha}{r} \rceil$ elements $a_{i,j}$ from all $k-1$ non-failed systematic nodes and $\lceil \frac{\alpha}{r} \rceil$ elements $p_{i,1}$ from p_1 where $i \in D_{\rho,d_l}$;
2: Repair $a_{i,l}$ where $i \in D_{\rho,d_l}$
3: Access and transfer $(r-1)\lceil \frac{\alpha}{r} \rceil$ elements $p_{i,j}$ from $p_2, \ldots, p_r$ where $i \in D_{\rho,d_l}$;
4: Access and transfer from the systematic nodes the elements $a_{i,j}$ listed in the $i-$th row of the arrays $P_2, \ldots, P_r$ where $i \in D_{\rho,d_l}$ that have not been read in Step 1;
5: Repair $a_{i,l}$ where $i \in \mathcal{D} \setminus D_{\rho,d_l}$;

Proposition 1: The repair bandwidth for a single systematic node γ is bounded between the following lower and upper bounds:

$$\frac{n-1}{r} \leq \gamma \leq \frac{(n-1)}{\alpha}\lceil \frac{\alpha}{r} \rceil + \frac{(r-1)}{\alpha}\lceil \frac{\alpha}{r} \rceil \lceil \frac{k}{r} \rceil. \quad (3)$$

Proof: Note that we read in total $k\lceil \frac{\alpha}{r} \rceil$ elements in Step 1 of Alg. 4. Additionally, $(r-1)\lceil \frac{\alpha}{r} \rceil$ elements are read in Step 3. Assuming that we do not read more elements in Step 4, we determine the lower bound as $k\lceil \frac{\alpha}{r} \rceil + (r-1)\lceil \frac{\alpha}{r} \rceil = (n-1)\lceil \frac{\alpha}{r} \rceil$ elements, i.e., the lower bound is $\frac{(n-1)}{\alpha}\lceil \frac{\alpha}{r} \rceil$ (since every element has a size of $\frac{1}{\alpha}$). To derive the upper bound, we assume that we read all elements $a_{i,j}$ from the extra $\lceil \frac{k}{r} \rceil$ columns of the arrays $P_2, \ldots, P_r$ in Step 4. Thus, the upper bound is $\frac{(n-1)}{\alpha}\lceil \frac{\alpha}{r} \rceil + \frac{(r-1)}{\alpha}\lceil \frac{\alpha}{r} \rceil \lceil \frac{k}{r} \rceil$. ■

The optimality of the proposed code construction is captured in the following Proposition.

Proposition 2: If r divides α, then the indexes (i,j) of the elements $a_{i,j}$ where $i \in \mathcal{D} \setminus D_{\rho,d_j}$ for each group of r systematic nodes are scheduled in one of the $\lceil \frac{k}{r} \rceil$ additional columns in the index arrays $P_2, \ldots, P_r$.

Next we show that there always exists a set of non-zero coefficients from $\mathbf{F}_q$ in the linear combinations so that the code is MDS. We adapt Theorem 4.1 from [6] as follows:

Theorem 1: There exists a choice of non-zero coefficients $c_{l,i,j}$ where $l = 1, \ldots, r$, $i = 1, \ldots, \alpha$ and $j = 1, \ldots, k$ from $\mathbf{F}_q$ such that the code is MDS if $q \geq \binom{n}{k} r\alpha$.

Proof: The system of linear equations in (2) defines a system of $r \times \alpha$ linear equations with $k \times \alpha$ variables. A repair of one failed node is given in Alg. 4, but for the sake of this proof, we explain the repair by discussing the solutions of the system of equations. When one node has failed, we have an overdefined system of $r \times \alpha$ linear equations with α unknowns. In general this can lead to a situation where there is no solution. However, since the values in system (2) are obtained from the values of the lost node, we know that there exists one solution. Thus, solving this system of $r \times \alpha$ linear equations with an overwhelming probability gives a unique solution, i.e., the lost node is recovered. When 2 nodes have failed, we have a system of $r \times \alpha$ linear equations with 2α unknowns. The same discussion for the overdefined system applies here. The most important case is when $r = n - k$ nodes have failed. In this case, we have a system of $r \times \alpha$ linear equations with $r \times \alpha$ unknowns. If the size of the finite field $\mathbf{F}_q$ is large enough, i.e., $q \geq \binom{n}{k} r\alpha$, as it is shown in Theorem 4.1 in [6], the system has a unique solution, i.e., the file M can be collected from any k nodes. ■

III. Code Examples

Let us take the $(5, 3, 4)$ code. We show a code construction for the optimal sub-packetization level $\alpha = 2^{\lceil \frac{3}{2} \rceil} = 4$.

The following requirements have to be satisfied for the code to be an access-optimal MDS code that achieves the lower bound of the repair bandwidth for any systematic node:

- $M = k\alpha = 12$ symbols,
- Repair a failed systematic node by accessing and transferring $\lceil \frac{\alpha}{r} \rceil = 2$ symbols from the remaining $d = 4$ nodes,
- Reconstruct the data from any 3 nodes.

The systematic nodes d_1, d_2, d_3 and the parity nodes p_1, p_2 are shown in Fig. 1. The file size is 12 symbols, where each node stores $\alpha = 4$ symbols. The elements of p_1 are linear

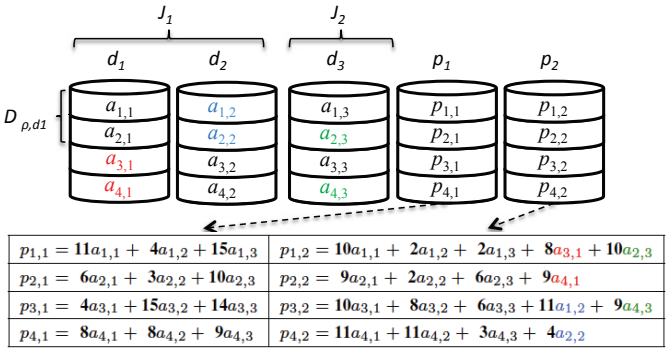

Fig. 1. An MDS array code with 3 systematic and 2 parity nodes for $\alpha = 4$. The elements presented in colors are scheduled as additional elements in p_2.The coefficients are from $\mathbf{F}_{16}$ with irreducible polynomial $x^4 + x^3 + 1$.

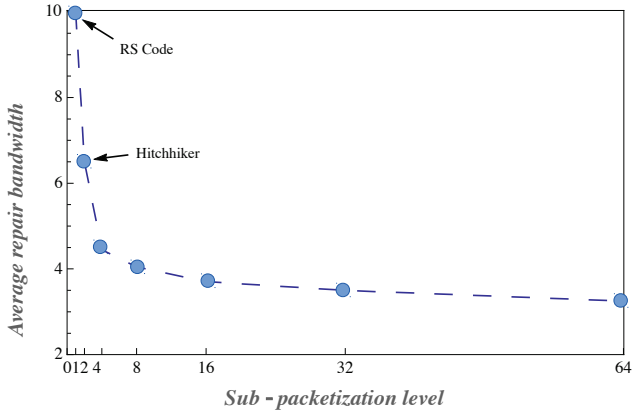

Fig. 2. Average repair bandwidth for any systematic node for different sub-packetization levels α for an (14, 10, 13) code

combinations of the row elements from the systematic nodes multiplied by coefficients from $\mathbf{F}_{16}$. The elements of p_2 are obtained by adding extra symbols to the row sum. We next show the scheduling of an element $a_{i,j}$ from a specific d_j where $i \in \mathcal{D} \setminus D_{\rho,d_j}$ at $portion = 2$ positions in the i-th row, $i \in D_{\rho,d_j}$, and the $(3+\nu)$-th column, $\nu = 1, 2$, of P_2. We follow the steps in Alg. 2 and give a brief description:

1. Initialize P_1 and P_2 as arrays $P = ((i,j))_{4\times 3}$.
2. Append additional 2 columns to P_2 initialized to $(0,0)$.
3. Set $portion = 2$ and $ValidPartitions = \emptyset$.
4. For the nodes d_1, d_2 that belong to J_1, run is equal to 2 and $step$ to 0. While run is equal to 1 and $step$ to 1 for the node d_3 that belongs to J_2.
5. Alg. 3 gives $D_{d_1} = \{\{1,2\},\{3,4\}\}$. Following step 14 in Alg. 2, the first 2 zero pairs in the 5-th column of P_2 with 0 distance between them are at the positions $(i,5)$ where $i = 1, 2$. Thus, $D_{\rho,d_1} = \{1,2\}$. We follow the same logic to obtain $D_{\rho,d_2} = \{3,4\}$ and $D_{\rho,d_3} = \{1,3\}$. Next we schedule the elements of d_j with i indexes that are not elements of D_{ρ,d_j} (represented in colors in Fig. 1) in the $i-$th row of P_2 where $i \in D_{\rho,d_j}$. The $i-$th index of $a_{3,1}$ and $a_{4,1}$ does not belong to D_{ρ,d_1} so these elements are scheduled at the positions $(1,5)$ and $(2,5)$ of P_2. We add the elements $a_{1,2}, a_{2,2}$ from d_2 in the 3-rd and the 4-th row respectively, while we add the elements $a_{2,3}$ and $a_{4,3}$ from d_3 in the 1-st and 3-rd row respectively. The symbols in the parity nodes are obtained as linear combinations of the row elements in the parity arrays. For instance, $p_{1,2}$ is a linear combination of the elements in the first row from all systematic nodes, $a_{3,1}$ and $a_{2,3}$.

We next show how to repair the node d_1 following Alg. 4. First, we repair the elements $a_{1,1}, a_{2,1}$. Thus, we access and transfer $a_{1,2}$, $a_{2,2}$, $a_{1,3}$ and $a_{2,3}$ from d_2 and d_3 and $p_{1,1}$, $p_{2,1}$ from p_1. In order to recover $a_{3,1}, a_{4,1}$, we need to access and transfer $p_{1,2}$ and $p_{2,2}$ from p_2. Hence, the data from d_1 is recovered by accessing and transferring in total 8 elements from 4 non-failed nodes. Exactly the same amount of data, 8 symbols, is needed to repair d_2 or d_3. Thus, the average repair bandwidth, defined as the ratio of the total repair bandwidth to repair all systematic nodes to the file size M, is equal to 2 symbols.

A. Performance Analysis

Another code discussed in this section is $(14, 10, 13)$ with different α. Fig. 2 shows the relation between the average repair bandwidth and α. For an RS code, α is 1 and the average repair bandwidth is equal to the file size. A Hitchhiker code [9] for $\alpha = 2$ reduces the repair bandwidth by 35% compared to the RS code. The remaining values of the average repair bandwidth are for the codes constructed with the algorithms presented in Section II. We observe that the lower bound of the repair bandwidth that is 3.25 is achieved for $\alpha = r^{\lceil \frac{k}{r} \rceil} = 64$. As we can see the repair bandwidth decreases as α increases.

IV. Conclusions

We presented a general construction of access-optimal regenerating codes that reach the lower bound of the repair bandwidth for $\alpha = r^{\lceil \frac{k}{r} \rceil}$, while the repair bandwidth is as close as possible to the lower bound when $\alpha < r^{\lceil \frac{k}{r} \rceil}$. The repair process of a systematic node is linear and highly parallelized.

References

[1] H. Weatherspoon and J. Kubiatowicz, "Erasure coding vs. replication: A quantitative comparison," in *Proc. 1st Int. Workshop on Peer-to-Peer Systems*, 2002, pp. 328–338.

[2] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4539–4551, Sept. 2010.

[3] I. Tamo, Z. Wang, and J. Bruck, "Zigzag codes: Mds array codes with optimal rebuilding," *IEEE Trans. Inf. Theory*, vol. 59, no. 3, pp. 1597–1616, March 2013.

[4] I. Tamo, Z. Wang, and J. Bruc, "Access vs. bandwidth in codes for storage," in *Proc. IEEE Int. Symp. Inf. Theory*, 2012, pp. 1187–1191.

[5] V. R. Cadambe, C. Huang, J. Li, and S. Mehrotra, "Polynomial length mds codes with optimal repair in distributed storage," in *Proc. 45th Asilomar Conf. Signals, Syst., Comp.*, 2011, pp. 1850–1854.

[6] G. K. Agarwal, B. Sasidharan, and P. V. Kumar, "An alternate construction of an access-optimal regenerating code with optimal sub-packetization level," in *Proc. 21st Nat. Conf. Comm.*, 2015, pp. 1–6.

[7] Z. Wang, I. Tamo, and J. Bruck, "On codes for optimal rebuilding access," in *Proc. 49th Annual Allerton Conf. Comm., Control, Comp.,*, 2011, pp. 1374–1381.

[8] B. Sasidharan, G. K. Agarwal, and P. V. Kumar, "A high-rate msr code with polynomial sub-packetization level," in *Proc. IEEE Int. Symp. Inf. Theory*, 2015, pp. 2051–2055.

[9] K. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, "A "hitchhiker's" guide to fast and efficient data reconstruction in erasure-coded data centers," in *Proc. ACM SIGCOMM Conf.*, 2014, pp. 331–342.